Algorithmic Game Theory and Applications: Homework 1

This homework is due at 3:00pm, on Monday, February 27th.
(This is a firm deadline. Please hand it in by that time to the ITO. Do not collaborate with other students on the coursework. Your solutions must be your own.) Each question counts for 20 points, for a total of 100 points.

1. Consider the following 2-player strategic game, G:

$$\begin{bmatrix} (6,8) & (2,9) & (3,8) & (2,8) \\ (0,5) & (2,3) & (2,6) & (8,4) \\ (7,0) & (2,7) & (4,4) & (4,3) \\ (2,3) & (5,3) & (2,5) & (5,4) \end{bmatrix}$$

   This is a "bimatrix", to be read as follows: Player 1 is the row player, and Player 2 is the column player. If the content of the bimatrix at row i and column j is the pair (a, b), then $u_1(i, j) = a$ and $u_2(i, j) = b$.

   Find (i.e., compute) *all* Nash equilibria for this game G.

   For any profile x that you claim is an NE of G, prove that x is indeed an NE of G.

   Argue why there are no other (pure or mixed) NEs, other than the profiles you claim are NEs.

2. Consider the 2-player zero-sum game given by the following payoff matrix for player 1 (the row player):

$$\begin{bmatrix} 7 & 1 & 6 & 4 & 2 \\ 2 & 6 & 8 & 3 & 5 \\ 6 & 3 & 8 & 3 & 4 \\ 5 & 5 & 4 & 7 & 4 \\ 2 & 8 & 2 & 8 & 9 \end{bmatrix}$$

   Set up a linear program associated with this game, and use some linear program solver to compute both the minimax value for this game, as well as a minimax profile, i.e., "optimal" (i.e., minmaximizer and maxminimizer) strategies for players 1 and 2, respectively. Specify the linear program(s) that you used to solve the game. Also, specify the *dual* of the linear program, and explain how to interpret the variables of the dual program.

(You can, for example, use the linear programming solver package `linprog` in MATLAB,available on DICE machines. To run MATLAB, type "matlab" at the shell command prompt. For guidance on using the `linprog` package, see:
`http://uk.mathworks.com/help/optim/ug/linprog.html`.)

3. Consider the following simple *2-player zero-sum* game called *Matching pennies*, where the payoff table for Player 1 (the row player) is given by:

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

We can view this as a game where each player chooses "heads" (H) or "tails" (T), and if their choices match, then player 1 wins, but if their choices don't match, then player 2 wins.

(a) (2 points) First, a very easy question: what is the unique Nash equilibrium, or equivalently the unique minimax profile of mixed strategies for the two players, in this game?

(b) (18 points) Now, suppose that the two players play this game against each other over and over again, for ever, and suppose that both of them use the following method in order to update their own (mixed) strategy after each round of the game.

i. At the beginning, in the first round, each player chooses either of the pure strategies, H or T, arbitrarily, and plays that.

ii. After each round, each player i accumulates statistics on how its opponents has played until now, meaning how many Heads and how many Tails have been played by the opponent, over all rounds of the game played thusfar. Suppose these numbers are N Heads and M Tails.
Then player i uses these statistics to "guess" its opponents *"statistical mixed strategy"* as follows. It assumes that its opponent will next play a mixed strategy σ, which plays Heads with probability $N/(N+M)$ and plays Tails with probability $M/(N+M)$.
Under the assumption that its opponent is playing the *"statistical mixed strategy"* σ, in the next round player i plays a pure strategy (H or T) that is a pure *best response* to σ.

If the statistical mixed strategy σ of the opponent is $(1/2, 1/2)$, then you are allowed to use **any tie breaking rule you wish** in order to determine the pure strategy played in the next round by player i.

iii. They repeat playing like this forever.

Do one of the following two things (preferably the first):

i. Prove that, regardless how the two players start playing the game in the first round, the "statistical mixed strategies" of both players in this method of repeatedly playing the matching pennies game will, in the long run, as the number of rounds goes to infinity, converge to their mixed strategies in the unique Nash equilibrium of the game.

   You are allowed to show that this holds using any specific tie breaking rule that you want. Please specify the precise tie breaking rule you have used. (It turns out that it hold true for any tie breaking rule. So it would be better, but not required, if you actually prove that any tie breaking rule works.)

ii. Alternatively, instead of proving that it works, you can "show" this **experimentally** by writing a simple program that plays this strategy update method for both players in repeated matching pennies, and show experimentally that, for all possible start strategies of both players, the "statistical mixed strategies" of the two players *looks like* it is converging to their NE strategies. (You will need to provide your program code, as well as the experimental output which shows that convergence "looks like" it is happening.)

   Note that experimentally you can only "show" that the "statistical mixed strategies" *look like* they are converging to the NE, by repeating the game some finite number of times, but you can not be sure that they do actually converge to the NE, without proving this. This is why a mathematical proof is preferable.

(c) (0 points. This is a really hard question that I want you to think about; but you do not need to submit a solution for it, unless you want to impress us, since you get zero marks for it.)

Does this same method of updating strategies always converge to *statistical mixed strategies* that yield a Nash Equilibrium for *any* finite 2-player normal form game? If so, explain why it does. If not, give an example of a 2-player finite game where it doesn't work, and argue why it doesn't work.

4. Consider the following 2-player strategic game, G:

$$\begin{bmatrix} (4,3) & (7,2) \\ (3,2) & (6,3) \end{bmatrix}$$

First, note that the pure strategy pair where both (row and column) players play their first pure strategy constitutes the only Nash Equilibrium in this game. In that NE, the payoff to player 1 (the row player) is 4.

Indeed, note that for the row player (player 1) the second row is strictly dominated by the first row.

Now imagine that player 1 is somehow able to *commit to* pure strategy 2 (the second row), in such a way that player 2 is convinced that this is what player 1 will do. Then player 2 must respond with pure strategy 2, which is *optimal against that strategy.* Note that in that case the payoff to player 1 will be 6 (and the payoff to player 2 will be 3). Therefore, player 1 has increased his own payoff by *committing to* a pure strategy that is strictly dominated.

(a) (3 points) Show that in this game the row player can increase its expected payoff even more by *committing to* a mixed strategy, rather than a pure strategy. In other words, show that there is a mixed strategy $x_1 = (x_{1,1}, x_{1,2})$ for player 1, such that if player 1 commits to x_1, and player 2 has to play optimally against x_1, then player 1 gets expected payoff strictly greater than 6.

(b) (3 points) Assume that whenever player 1 *commits to* a mixed strategy x_1, and player 2 has more than one optimal counter-strategy, player 2 plays a pure counter-strategy which maximizes its own expected payoff, and among all pure strategies that do, maximizes player 1's expected payoff. (In other words, whenever player 2 is indifferent between different pure strategies to play against x_1, player 2 chooses one such optimal pure strategy that is also the most beneficial one for player 1.)

Under these rules for playing, what is the *optimal mixed strategy*, x_1^*, that player 1 can *commit to* in the above game, in order to maximize its own expected payoff, assuming player 2 will play against it in the above fashion?

(c) (14 points)

Now consider a general finite 2-player strategic form game, with m pure strategies for player 1, and n pure strategies for player 2. The game is specified by two $m \times n$ matrices, A and B, which specify the payoff matrix for player 1 and player 2, respectively.

Give an algorithm for computing an optimal mixed strategy, x_1^*, that player 1 can *commit to*, assuming player 2 plays against such a strategy in the manner specified above.

Explain why your algorithm is correct.

(Hint: for each pure strategy s of player 2, consider separately the case where s is a best response to player 1's commitment strategy x_1. In each such case, set up a new linear program (LP) to compute the best mixed strategy x_1 for player 1, under the constraint that s is a best response against it. Then combine the results of these LPs to obtain the best overall mixed strategy to commit to for player 1.)

5. (This is more challenging than the prior problems. But still doable, especially given the hint.)

One variant of the Farkas Lemma says the following:

Farkas Lemma A linear system of inequalities $Ax \leq b$ has a solution x if and only if there is no vector y satisfying $y \geq 0$ and $y^T A = 0$ (i.e., 0 in every coordinate) and such that $y^T b < 0$.

Prove this Farkas Lemma with the aid of Fourier-Motzkin elimination. (*Hint:* One direction of the "if and only if" is easy. For the other direction, use induction on the number of columns of A, using the fact that Fourier-Motzkin elimination "works". Note basically that each round of Fourier-Motzkin elimination can "eliminate one variable" by pre-multiplying a given system of linear inequalities by a *non-negative* matrix.)